arXiv:1408.6323v2 [math.ST] 3 Sep 2014

ON BAYESIAN A- AND D-OPTIMAL EXPERIMENTAL DESIGNS IN INFINITE DIMENSIONS

ALEN ALEXANDERIAN, PHILIP GLOOR, AND OMAR GHATTAS

ABSTRACT. We consider Bayesian linear inverse problems in infinite-dimensional separable Hilbert spaces, with a Gaussian prior measure and additive Gaussian noise model, and provide an extension of the concept of Bayesian D-optimality to the infinite-dimensional case. To this end, we derive the infinite-dimensional version of the expression for the Kullback-Leibler divergence from the posterior measure to the prior measure, which is subsequently used to derive the expression for the expected information gain. We also study the notion of Bayesian A-optimality in the infinite-dimensional setting, and extend the well known (in the finite-dimensional case) equivalence of the Bayes risk of the MAP estimator with the trace of the posterior covariance, for the Gaussian linear case, to the infinite-dimensional Hilbert space case.

1. INTRODUCTION

In a Bayesian inference problem one uses experimental (observed) data to update the prior state of knowledge about a parameter which often specifies certain properties of a mathematical model. The ingredients of a Bayesian inference problem include the prior measure which encodes our prior knowledge about the inference parameter, experimental data, and the data likelihood which describes the conditional distribution of the experimental data for a given model parameter. The solution of a Bayesian inference problem is a posterior probability law for the inference parameter. The quality of this solution, which can be measured using different criteria, depends to a large extent on the experimental data used in solving the inference problem. In practice, acquisition of such experimental data is often costly, as it requires deployment of scarce resources. Hence, the problem of optimal collection of experimental data, i.e. that of optimal experimental design (OED) [2, 17, 13], is an integral part of modeling under uncertainty. The basic problem of OED is to optimize a function of the experimental setup which describes, in a certain sense which needs to be specified, the statistical quality of the solution to the Bayesian inference problem. Note that what constitutes an experimental design depends on the application at hand. For example, in a problem involving diffusive transport of a contaminant, one may use measurements of concentration at sensor sites in the physical domain (at a certain point in time) to *infer* where the contaminant originated, i.e. the initial state of the concentration field. In this problem, an experimental design specifies the locations of the sensors in the physical domain. Note also that the inference parameter in this example, i.e. the initial concentration field, is a random function (random field) whose realizations belong to an appropriate function space.

Date: September 4, 2014.

2010 *Mathematics Subject Classification.* 62K05;62F15;46N30;49N45.

Key words and phrases. Bayesian inference in Hilbert space; Gaussian measure; Kullback Leibler divergence; Bayesian optimal experimental design; expected information gain; Bayes risk.

We consider the problem of design of experiments for inference problems whose inference parameter belongs to an infinite-dimensional separable Hilbert space. This is motivated by the recent interest in the Bayesian framework for inverse problems [16]. A *Bayesian inverse problem* involves inference of Hilbert space valued parameters that describe physical properties of mathematical models which are often governed by partial differential equations. Study of such problems requires a synthesis of ideas from inverse problem theory, PDE-constrained optimization, functional analysis, and probability and statistics and has provided a host of interesting mathematical problems with a wide range of applications. The problem of design of experiments in this infinite-dimensional setting involves optimizing functionals of experimental designs which are defined in terms of operators on Hilbert spaces.

The precise definition of what is meant by an *optimal* design leads to the choice of a design criterion. A popular experimental design criterion, in the finite-dimensional case, is that of D-optimality which seeks to minimize the determinant of the posterior covariance operator. The geometric intuition behind D-optimality is that of minimizing the *volume* of the uncertainty ellipsoid. Minimizing this determinant, however, is not meaningful in infinite dimensions, as the posterior covariance operator is a trace-class linear operator with eigenvalues that accumulate at zero. In the present work, we provide an extension of the concept of D-optimal design to the infinite-dimensional Hilbert space setting. In particular, we focus on the case of Bayesian linear inverse problems whose parameter space is an infinite-dimensional separable Hilbert space which we denote by $\mathscr{H}$, and we assume a Gaussian prior measure, and an additive Gaussian noise model. To study the concept of D-optimality in the infinite-dimensional setting we formulate the problem as that of maximizing the expected information gain, measured by the Kullback-Leibler (KL) divergence [12] from posterior to prior. To be precise, if μ_{pr} denotes the prior measure, $\boldsymbol{y}$ is a vector of experimental data obtained using an experimental design specified by a vector of design parameters $\boldsymbol{\xi}$, and $\mu_{\mathsf{post}}^{\boldsymbol{y},\boldsymbol{\xi}}$ denotes the resulting posterior measure, the KL divergence from posterior to prior is given by,

$$D_{\mathsf{kl}}\left(\mu_{\mathsf{post}}^{\boldsymbol{y},\boldsymbol{\xi}} \| \mu_{\mathsf{pr}}\right) := \int_{\mathscr{H}} \log\left\{\frac{d\mu_{\mathsf{post}}^{\boldsymbol{y},\boldsymbol{\xi}}}{d\mu_{\mathsf{pr}}}\right\} d\mu_{\mathsf{post}}^{\boldsymbol{y},\boldsymbol{\xi}}.$$

(The argument of the logarithm in the above formula is the Radon-Nikodym derivative of the posterior measure with respect to the prior measure.) The experimental design criterion is then defined by averaging $D_{\mathsf{kl}}(\mu_{\mathsf{post}}^{\boldsymbol{y},\boldsymbol{\xi}} \| \mu_{\mathsf{pr}})$ over all possible experimental data. In a Bayesian inverse problem, this averaging over experimental data can be done as follows:

$$\text{expected information gain} := \int_{\mathscr{H}} \int_{\mathscr{Y}} D_{\mathsf{kl}}\left(\mu_{\mathsf{post}}^{\boldsymbol{y},\boldsymbol{\xi}} \| \mu_{\mathsf{pr}}\right)\, \pi_{\mathsf{like}}(\boldsymbol{y}|u;\boldsymbol{\xi}) d\boldsymbol{y}\, \mu_{\mathsf{pr}}(du),$$

where $\boldsymbol{\xi}$ is a fixed design vector, $\mathscr{Y}$ denotes the space of experimental data and $\pi_{\mathsf{like}}(\boldsymbol{y}|u;\boldsymbol{\xi})$ is the data likelihood which specifies the distribution of $\boldsymbol{y}$ for a given $u \in \mathscr{H}$.

It is known in the finite-dimensional Gaussian linear case (i.e., an inference problem with Gaussian prior and noise distributions) that maximizing this expected information gain is equivalent to minimizing the determinant of the posterior covariance operator, i.e., the usual D-optimal design problem. While this does not directly extend to the infinite-dimensional case, it suggests a mathematically rigorous path to an infinite-dimensional analogue of Bayesian D-optimality. In the present work, we derive analytic expressions for the KL divergence from posterior to prior in a Hilbert space. This

enables deriving the expression for the expected information gain, leading to the infinite-dimensional version of the Bayesian D-optimal experimental design criterion.

We also discuss another popular experimental design criterion, that of A-optimality, in the infinite-dimensional setting. An A-optimal design is one that minimizes the trace of the posterior covariance operator; i.e., if $\mathcal{C}_{\text{post}}(\boldsymbol{\xi}) : \mathscr{H} \to \mathscr{H}$ denotes the posterior covariance operator corresponding to an experimental design $\boldsymbol{\xi}$, we seek to minimize $\operatorname{tr}(\mathcal{C}_{\text{post}}(\boldsymbol{\xi}))$. In the statistics literature it is known (see e.g., [6]) that for a Gaussian linear inference problem in $\mathscr{H} = \mathbb{R}^n$, minimizing the trace of the posterior covariance *matrix* is equivalent to minimizing the average mean square error of the maximum a posteriori probability (MAP) estimator for the inference parameter. We provide an extension of this result to the infinite-dimensional Hilbert space setting, where we show that the trace of the posterior covariance *operator*—a positive, self-adjoint, and trace-class operator on $\mathscr{H}$—coincides with the average mean square error of the MAP estimator.

Note that the design vector $\boldsymbol{\xi}$ enters the Bayesian inverse problem through the data likelihood. The exact nature of this dependence on $\boldsymbol{\xi}$ is not essential to our discussion and hence, to keep the notation simple, we suppress the dependence to $\boldsymbol{\xi}$ in our derivations. (See e.g., [6] for a an overview of how an experimental design is incorporated in an inference problem in classical formulations.)

2. Background concepts

In this section, we outline the background concepts that are needed in the rest of this article. In what follows, $\mathscr{H}$ denotes an infinite-dimensional separable real Hilbert space, with inner-product $\langle \cdot, \cdot \rangle_{\mathscr{H}}$ and induced norm $\|\cdot\|_{\mathscr{H}} = \langle \cdot, \cdot \rangle_{\mathscr{H}}^{1/2}$.

2.1. Trace-class operators on $\mathscr{H}$.

Let $\mathscr{L}(\mathscr{H})$ denote the set of bounded linear operators on $\mathscr{H}$. We say $\mathcal{A} \in \mathscr{L}(\mathscr{H})$ is positive if $\langle x, \mathcal{A}x \rangle_{\mathscr{H}} \geq 0$ for all $x \in \mathscr{H}$, and is strictly positive if $\langle x, \mathcal{A}x \rangle_{\mathscr{H}} > 0$ for all non-zero $x \in \mathscr{H}$. For $\mathcal{A} \in \mathscr{L}(\mathscr{H})$, $|\mathcal{A}| = (\mathcal{A}^*\mathcal{A})^{1/2}$, where $\mathcal{A}^*$ denotes the adjoint of $\mathcal{A}$. We say $\mathcal{A}$ is of *trace-class* if for any orthonormal basis $\{f_j\}_{j=1}^{\infty}$ of $\mathscr{H}$,

$$\sum_{j=1}^{\infty} \langle |\mathcal{A}| f_j, f_j \rangle_{\mathscr{H}} < \infty.$$

It is straightforward to show that the value of the above summation is invariant with respect to the choice of the orthonormal basis [14]. We denote by $\mathscr{L}_1(\mathscr{H})$ the subspace of $\mathscr{L}(\mathscr{H})$ consisting of trace-class operators. For $\mathcal{A} \in \mathscr{L}_1(\mathscr{H})$,

$$\operatorname{tr}(\mathcal{A}) = \sum_{j=1}^{\infty} \langle \mathcal{A} f_j, f_j \rangle_{\mathscr{H}},$$

where the sum is finite and its value is independent of the choice of the orthonormal basis [7, 14].

Let $\mathscr{L}_1^{\text{sym}+}(\mathscr{H})$ be the subspace of positive self-adjoint operators in $\mathscr{L}_1(\mathscr{H})$, and note that for $\mathcal{A} \in \mathscr{L}_1^{\text{sym}+}(\mathscr{H})$, there exists an orthonormal basis of eigenvectors, $\{e_j\}$, with corresponding (real, non-negative) eigenvalues, $\{\lambda_j\}$, and $\operatorname{tr}(A) = \sum_{j=1}^{\infty} \langle \mathcal{A} e_j, e_j \rangle_{\mathscr{H}} = \sum_{j=1}^{\infty} \lambda_j$.

In what follows we shall make repeated use of the following result: if $\mathcal{A} \in \mathscr{L}_1(\mathscr{H})$ and $\mathcal{B} \in \mathscr{L}(\mathscr{H})$ then $\mathcal{A}\mathcal{B}$ and $\mathcal{B}\mathcal{A}$ both belong to $\mathscr{L}_1(\mathscr{H})$ and $\operatorname{tr}(\mathcal{A}\mathcal{B}) = \operatorname{tr}(\mathcal{B}\mathcal{A})$; see e.g., [14] for details. Moreover, it is straightforward to show that if $\mathcal{A}$ is a trace-class

operator and $\mathcal{B} : \mathscr{H} \to \mathbb{R}^q$ is a bounded linear operator, then $\mathcal{A}\mathcal{B}^*\mathcal{B} \in \mathscr{L}_1(\mathscr{H})$ and $\operatorname{tr}(\mathcal{A}\mathcal{B}^*\mathcal{B}) = \operatorname{tr}(\mathcal{B}\mathcal{A}\mathcal{B}^*)$.

2.2. **Borel probability measures on $\mathscr{H}$.** We work with probability measures on the measurable space $\big(\mathscr{H}, \mathfrak{B}(\mathscr{H})\big)$, where $\mathfrak{B}(\mathscr{H})$ denotes the Borel sigma-algebra on $\mathscr{H}$; we refer to such measures as Borel probability measures. Let μ be a Borel probability measure on $\mathscr{H}$, which has bounded first and second moments. The mean $m \in \mathscr{H}$ and covariance operator $\mathcal{Q} \in \mathscr{L}(\mathscr{H})$ of μ are characterized as follows:

$$\langle m, x\rangle_{\mathscr{H}} = \int_{\mathscr{H}} \langle z, x\rangle_{\mathscr{H}}\ \mu(dz), \qquad \langle \mathcal{Q}x, y\rangle_{\mathscr{H}} = \int_{\mathscr{H}} \langle x, z-m\rangle_{\mathscr{H}}\, \langle y, z-m\rangle_{\mathscr{H}}\ \mu(dz),$$

for all $x, y \in \mathscr{H}$. It is straightforward to show (see e.g., [8]) that $\mathcal{Q}$ belongs to $\mathscr{L}_1^{\mathsf{sym}+}(\mathscr{H})$, and that

$$\int_{\mathscr{H}} \|x\|_{\mathscr{H}}^2\ \mu(dx) = \operatorname{tr}(\mathcal{Q}) + \|m\|_{\mathscr{H}}^2\,. \tag{1}$$

2.3. **Gaussian measures on $\mathscr{H}$.** In the present work, we shall be working with Gaussian measures on Hilbert spaces [8]; μ is a Gaussian measure on $(\mathscr{H}, \mathfrak{B}(\mathscr{H}))$ if for every $x \in \mathscr{H}$ the linear functional $\langle x, \cdot\rangle_{\mathscr{H}}$, considered as a random variable from $(\mathscr{H}, \mathfrak{B}(\mathscr{H}), \mu)$ to $(\mathbb{R}, \mathfrak{B}(\mathbb{R}))$, is a (one-dimensional) Gaussian random variable. We refer the reader to [8] or [9] for the theory of Gaussian measures on Hilbert spaces. We denote a Gaussian measure with mean $m \in \mathscr{H}$ and $\mathcal{Q} \in \mathscr{L}_1^{\mathsf{sym}+}(\mathscr{H})$ by $\mathcal{N}(m, \mathcal{Q})$. If $\mathcal{Q}$ satisfies $\ker(\mathcal{Q}) = \{0\}$, where $\ker(\mathcal{Q})$ denotes the null space of $\mathcal{Q}$, we say that $\mathcal{N}(m, \mathcal{Q})$ is a non-degenerate Gaussian measure.

In what follows, we shall use the following result, concerning the law of an affine transformation on $\mathscr{H}$: If $\mu = \mathcal{N}(m, \mathcal{Q})$, a Gaussian measure, $\mathcal{A} \in \mathscr{L}(\mathscr{H})$, and $b \in \mathscr{H}$, then $Tx = \mathcal{A}x + b$ is a random variable on $\mathscr{H}$ whose law is given by $\mu_T = \mu \circ T^{-1} = \mathcal{N}(\mathcal{A}m + b, \mathcal{A}\mathcal{Q}\mathcal{A}^*)$ [8]. Thus, in particular, we note that,

$$\int_{\mathscr{H}} \|Tx\|_{\mathscr{H}}^2\ \mu(dx) = \int_{\mathscr{H}} \|\xi\|_{\mathscr{H}}^2\ \mu_T(d\xi) = \operatorname{tr}(\mathcal{A}\mathcal{Q}\mathcal{A}^*) + \|\mathcal{A}m + b\|_{\mathscr{H}}^2\,,$$

where the last equality uses (1). It follows that if $\mathcal{A} \in \mathscr{L}(\mathscr{H})$ is positive self-adjoint compact operator, and $\mu = \mathcal{N}(m, \mathcal{Q})$ is a Gaussian measure, then

$$\begin{aligned}\int_{\mathscr{H}} \langle \mathcal{A}x, x\rangle_{\mathscr{H}}\ \mu(dx) &= \int_{\mathscr{H}} \left\|\mathcal{A}^{1/2}x\right\|_{\mathscr{H}}^2\ \mu(dx)\\ &= \operatorname{tr}(\mathcal{A}^{1/2}\mathcal{Q}\mathcal{A}^{1/2}) + \left\langle \mathcal{A}^{1/2}m, \mathcal{A}^{1/2}m\right\rangle_{\mathscr{H}}\\ &= \operatorname{tr}(\mathcal{A}\mathcal{Q}) + \langle \mathcal{A}m, m\rangle_{\mathscr{H}}\,.\end{aligned} \tag{2}$$

This shows that the well-known expression for the expectation of a quadratic form on $\mathbb{R}^n$ extends to the infinite-dimensional Hilbert space setting. It can be shown that, as in the finite-dimensional case, this result holds not just for Gaussian measures, but also for any Borel probability measure with mean m and covariance operator $\mathcal{Q}$; moreover, the the only requirement on the operator $\mathcal{A}$ is boundedness. That is, we have the following result:

Lemma 1. *Let μ be a Borel probability measure on $\mathscr{H}$ with mean $m \in \mathscr{H}$ and covariance operator $\mathcal{Q} \in \mathscr{L}_1^{\mathsf{sym}+}(\mathscr{H})$, and let $\mathcal{A} \in \mathscr{L}(\mathscr{H})$. Then,*

$$\int_{\mathscr{H}} \langle \mathcal{A}x, x\rangle_{\mathscr{H}}\ \mu(dx) = \operatorname{tr}(\mathcal{A}\mathcal{Q}) + \langle \mathcal{A}m, m\rangle_{\mathscr{H}}\,.$$

Proof. See Appendix A. □

2.4. **Kullback-Leibler divergence.** In probability theory the Kullback-Leibler (KL) divergence, also referred to as the relative entropy, is a measure of "distance" between two probability measures. This notion was defined in [12]. While KL divergence is not a metric—it is non-symmetric and does not satisfy the triangle inequality—it is used commonly in probability theory to describe the distance of a measure μ from a reference measure μ_0. Also, KL divergence does satisfy some of the intuitive notions of distance; i.e. the KL divergence from μ to μ_0 is non-negative and is zero if and only if the two measures are the same. Consider μ and μ_0 be two Borel probability measures and suppose μ is absolutely continuous with respect to μ_0. The KL divergence from μ to μ_0, denoted by $D_{\mathsf{kl}}(\mu\|\mu_0)$, is defined as

$$D_{\mathsf{kl}}(\mu\|\mu_0) = \int_{\mathscr{H}} \log\left\{\frac{d\mu}{d\mu_0}\right\} d\mu,$$

where $\frac{d\mu}{d\mu_0}$ is the Radon-Nikodym derivative of μ with respect to μ_0. In the case μ is not absolutely continuous with respect to μ_0 the KL divergence is $+\infty$. Notice that for Borel probability measures on $\mathbb{R}^n$ that admit densities with respect to the Lebesgue measure, we may rewrite the definition of the KL divergence in terms the densities; that is, if p and p_0 are Lebesgue densities, i.e., probability density functions (pdfs), of μ and μ_0 respectively, one has $D_{\mathsf{kl}}(\mu\|\mu_0) = \int_{\mathbb{R}^n} \log\big(p(\boldsymbol{x})/p_0(\boldsymbol{x})\big)\, p(\boldsymbol{x})\, d\boldsymbol{x}$. However, in an infinite-dimensional Hilbert space, where there is no Lebesgue measure, we are forced to work with the abstract definition of KL divergence presented above.

In this paper, we will be dealing with (non-degenerate) Gaussian measures on infinite-dimensional Hilbert spaces. For Gaussian measures on $\mathbb{R}^n$, one can use the expression for the (multivariate) Gaussian pdfs to derive the well-known analytic expression for the KL divergence between Gaussians. In the infinite-dimensional Hilbert space setting, not only do we not have access to pdfs, but given two Gaussian measures they are not necessarily equivalent.[1] In fact, given a centered Gaussian measure μ on $\mathscr{H}$, shifting the mean gives, μ-almost surely, a Gaussian measure which is singular with respect to μ; see e.g., [8, Chapter 2]. However, In the present work, we work with a special case, namely that of a Bayesian linear inverse problem on $\mathscr{H}$ with a Gaussian prior and an additive Gaussian noise model; in this case the posterior measure is also Gaussian and is equivalent to the prior [16], and thus, $D_{\mathsf{kl}}\left(\mu_{\mathsf{post}}^{\boldsymbol{y}}\|\mu_{\mathsf{pr}}\right)$ is well-defined. Later in the paper, we will derive the expression for the KL divergence from posterior to prior in an infinite-dimensional Hilbert space, which we shall use to derive the expression for the expected information gain.

3. Bayesian linear inverse problems in a Hilbert space

We consider the problem of inference of a parameter u which belongs to an infinite-dimensional Hilbert space $\mathscr{H}$. All our prior knowledge regarding the parameter u is encoded in a Borel probably measure on $\mathscr{H}$, which we refer to as the prior measure and denote by μ_{pr}; here we assume that μ_{pr} is a Gaussian measure $\mu_{\mathsf{pr}} = \mathcal{N}(u_{\mathsf{pr}}, \mathcal{C}_{\mathsf{pr}})$. Moreover, in what follows, we assume that $\ker(\mathcal{C}_{\mathsf{pr}}) = \{0\}$, i.e., μ_{pr} is non-degenerate. The inference problem uses experimental data $\boldsymbol{y} \in \mathscr{Y}$ to update the prior state of knowledge on the law of the parameter u. Here $\mathscr{Y}$ is the space of the experimental data, which in the present work is $\mathscr{Y} = \mathbb{R}^q$. We assume that u is a model parameter which is related to experimental data $\boldsymbol{y} \in \mathscr{Y}$ according to the following noise model,

$$\boldsymbol{y} = \mathcal{G}u + \boldsymbol{\eta}. \tag{3}$$

[1] Recall that two measures are called equivalent if they are mutually absolutely continuous with respect to each other.

The operator $\mathcal{G} : \mathscr{H} \to \mathscr{Y}$ is the *parameter-to-observable map* and is assumed to be a continuous linear mapping. In practice, for a given u, computing $\mathcal{G}u$ would involve the evaluation of a mathematical model with the parameter value u followed by the application of a restriction operator to extract data at pre-specified locations in space and/or time. The discrepancy between the model output $\mathcal{G}u$ and experimental data $\boldsymbol{y}$ is modeled by $\boldsymbol{\eta}$ which is a random vector that accounts for experimental noise, i.e. noise associated with the process of collecting experimental data. We assume $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Gamma}_{\text{noise}})$, and thus, the distribution of $\boldsymbol{y}|u$ is Gaussian, $\boldsymbol{y}|u \sim \mathcal{N}(\mathcal{G}u, \mathbf{\Gamma}_{\text{noise}})$ with pdf

$$\pi_{\text{like}}(\boldsymbol{y}|u) = \frac{1}{\mathcal{Z}_{\text{like}}} \exp\Big\{ -\frac{1}{2}(\mathcal{G}u - \boldsymbol{y})^T \mathbf{\Gamma}_{\text{noise}}^{-1}(\mathcal{G}u - \boldsymbol{y})\Big\},$$

where $\mathcal{Z}_{\text{like}} = (2\pi)^{q/2} \det(\mathbf{\Gamma}_{\text{noise}})^{1/2}$. In what follows, we denote

$$\Phi(u; \boldsymbol{y}) = \frac{1}{2}(\mathcal{G}u - \boldsymbol{y})^T \mathbf{\Gamma}_{\text{noise}}^{-1}(\mathcal{G}u - \boldsymbol{y}). \tag{4}$$

3.1. **The Bayes formula and the posterior measure.** The solution of the Bayesian inverse problem is the posterior measure, describing the law of the parameter u, conditioned on the experimental data $\boldsymbol{y}$, and is linked to the prior measure μ_{pr} through the infinite-dimensional version of Bayes Theorem [16]:

$$\frac{d\mu_{\text{post}}^{\boldsymbol{y}}}{d\mu_{\text{pr}}} = \frac{1}{\mathcal{Z}(\boldsymbol{y})} \pi_{\text{like}}(\boldsymbol{y}|u), \tag{5}$$

where $\mathcal{Z}(\boldsymbol{y})$ is the normalization constant. Notice that we can rewrite Bayes Theorem as,

$$\frac{d\mu_{\text{post}}^{\boldsymbol{y}}}{d\mu_{\text{pr}}} = \frac{1}{\mathcal{Z}_0(\boldsymbol{y})} \exp\{-\Phi(u; \boldsymbol{y})\}, \tag{6}$$

with $\mathcal{Z}_0(\boldsymbol{y}) = \int_{\mathscr{H}} \exp\{-\Phi(u; \boldsymbol{y})\}\, \mu_{\text{pr}}(du)$. In the Gaussian linear case, it is possible to evaluate $\mathcal{Z}_0(\boldsymbol{y})$ analytically; see Lemma 2 below.

As discussed above, we consider Bayesian linear inverse problems; i.e., Bayesian inverse problems involving a linear parameter-to-observable map $\mathcal{G}$. It is well known [16] that for a Gaussian linear inverse problem, as specified above, the solution is a Gaussian posterior measure $\mu_{\text{post}}^{\boldsymbol{y}} = \mathcal{N}\big(u_{\text{post}}^{\boldsymbol{y}}, \mathcal{C}_{\text{post}}\big)$ with,

$$\mathcal{C}_{\text{post}} = (\mathcal{G}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \mathcal{G} + \mathcal{C}_{\text{pr}}^{-1})^{-1}, \qquad u_{\text{post}}^{\boldsymbol{y}} = \mathcal{C}_{\text{post}}(\mathcal{G}^* \mathbf{\Gamma}_{\text{noise}}^{-1} \boldsymbol{y} + \mathcal{C}_{\text{pr}}^{-1} u_{\text{pr}}).$$

In practice, the noise covariance matrix, $\mathbf{\Gamma}_{\text{noise}}$ is often a multiple of the identity, $\mathbf{\Gamma}_{\text{noise}} = \sigma^2 \mathbf{I}$, where σ is the experimental noise level. In the derivations that follow, since there is no loss of generality, we take $\sigma = 1$. Generalizing the results to the cases where $\mathbf{\Gamma}_{\text{noise}}$ is an anisotropic diagonal matrix (uncorrelated observations with varying experimental noise levels) or more generally $\mathbf{\Gamma}_{\text{noise}}$ that is symmetric and positive definite with nonzero off diagonal entries (correlated observations) is straightforward. Moreover, for simplicity, we assume that the prior is a centered Gaussian, i.e. $u_{\text{pr}} = 0$. Again, the generalization to the case of non-centered prior measure is straightforward. With these simplifications, the mean and covariance of the posterior measure are given by,

$$\mathcal{C}_{\text{post}} = (\mathcal{G}^* \mathcal{G} + \mathcal{C}_{\text{pr}}^{-1})^{-1}, \qquad u_{\text{post}}^{\boldsymbol{y}} = \mathcal{C}_{\text{post}} \mathcal{G}^* \boldsymbol{y}. \tag{7}$$

In what follows, we use the notation,

$$\mathcal{H}_{\text{m}} = \mathcal{G}^* \mathcal{G}. \tag{8}$$

The motivation behind this notation is that $\mathcal{G}^*\mathcal{G}$ is the Hessian of the functional, $\Phi(u; \boldsymbol{y})$, which measures the magnitude of the *misfit* between experimental data $\boldsymbol{y}$ and

model prediction $\mathcal{G}u$. Note that in statistical terms, $\mathcal{H}_{\mathsf{m}}$ is the Hessian of the negative log-likelihood which is also referred to as the *Fisher information matrix*. Another notation we shall use frequently is,

$$\tilde{\mathcal{H}}_{\mathsf{m}} = \mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}^{1/2}. \tag{9}$$

Intuitively, this *prior-preconditioned* $\mathcal{H}_{\mathsf{m}}$ can be thought of as the information matrix which has been filtered through the prior. To further appreciate the notion of the prior-preconditioned misfit Hessian, we note that the second moment of the parameter-to-observable map, considered as a random variable $\mathcal{G} : (\mathscr{H}, \mathfrak{B}(\mathscr{H}), \mu_{\mathsf{pr}}) \to (\mathbb{R}^q, \mathfrak{B}(\mathbb{R}^q))$ is given by,

$$\int_{\mathscr{H}} |\mathcal{G}u|^2 \, \mu_{\mathsf{pr}}(du) = \int_{\mathscr{H}} \langle \mathcal{G}u, \mathcal{G}u \rangle_{\mathbb{R}^q} \, \mu_{\mathsf{pr}}(du) = \int_{\mathscr{H}} \langle \mathcal{H}_{\mathsf{m}}u, u \rangle_{\mathscr{H}} \, \mu_{\mathsf{pr}}(du) = \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}\mathcal{H}_{\mathsf{m}}) = \mathrm{tr}(\tilde{\mathcal{H}}_{\mathsf{m}}).$$

3.2. **A spectral point of view of uncertainty reduction.** Let $\tilde{\mathcal{H}}_{\mathsf{m}}$ be the prior-preconditioned misfit Hessian as defined in (9) and denote

$$\mathcal{S} = (I + \tilde{\mathcal{H}}_{\mathsf{m}})^{-1}. \tag{10}$$

The posterior covariance operator, $\mathcal{C}_{\mathsf{post}}$, given in (7) can be written as, $\mathcal{C}_{\mathsf{post}} = \mathcal{C}_{\mathsf{pr}}^{1/2}(I + \tilde{\mathcal{H}}_{\mathsf{m}})^{-1}\mathcal{C}_{\mathsf{pr}}^{1/2} = \mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{S}\mathcal{C}_{\mathsf{pr}}^{1/2}$. We consider the quantity,

$$\delta(\mathcal{C}_{\mathsf{pr}}, \mathcal{C}_{\mathsf{post}}) := \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}) - \mathrm{tr}(\mathcal{C}_{\mathsf{post}}) = \mathrm{tr}\big(\mathcal{C}_{\mathsf{pr}}^{1/2}(I - \mathcal{S})\mathcal{C}_{\mathsf{pr}}^{1/2}\big).$$

For the class of Bayesian linear inverse problems considered in the present work, it is straightforward to show that $\delta(\mathcal{C}_{\mathsf{pr}}, \mathcal{C}_{\mathsf{post}}) \geq 0$. In particular, we note that if $\{\lambda_i\}$ and $\{e_i\}$ are the eigenvalues and the respective eigenvectors of $\tilde{\mathcal{H}}_{\mathsf{m}}$, then

$$\langle e_i, (I - \mathcal{S})e_i \rangle_{\mathscr{H}} = 1 - \langle e_i, \mathcal{S}e_i \rangle_{\mathscr{H}} = 1 - 1/(1 + \lambda_i) = \lambda_i/(1 + \lambda_i) \geq 0, \quad i = 1, 2, \ldots,$$

which shows that $\delta(\mathcal{C}_{\mathsf{pr}}, \mathcal{C}_{\mathsf{post}}) = \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}^{1/2}(I - \mathcal{S})\mathcal{C}_{\mathsf{pr}}^{1/2}) \geq 0$. The quantity $\delta(\mathcal{C}_{\mathsf{pr}}, \mathcal{C}_{\mathsf{post}})$ can thus be considered a measure of variance (uncertainty) reduction. More precisely, we consider for each $i \geq 1$,

$$\langle e_i, \mathcal{C}_{\mathsf{post}}e_i \rangle_{\mathscr{H}} = \int_{\mathscr{H}} \big\langle e_i, u - u_{\mathsf{post}}^{\boldsymbol{y}} \big\rangle_{\mathscr{H}}^2 \, \mu_{\mathsf{post}}^{\boldsymbol{y}}(du),$$

which measures the posterior variance of the coordinate of u in the direction e_i.

Proposition 1. *Let $\{\lambda_i, e_i\}_1^\infty$ be eigenpairs of $\tilde{\mathcal{H}}_m$. Then, $\langle e_i, \mathcal{C}_{post}e_i \rangle_{\mathscr{H}} \leq \langle e_i, \mathcal{C}_{pr}e_i \rangle_{\mathscr{H}}$, for all $i \geq 1$.*

Proof. Note that for each $v \in \mathscr{H}$, $\mathcal{S}v = \sum_j (1 + \lambda_j)^{-1} \langle e_j, v \rangle_{\mathscr{H}} \, e_j$. Hence,

$$\begin{aligned}
&\langle e_i, \mathcal{C}_{\mathsf{post}}e_i \rangle_{\mathscr{H}} = \Big\langle e_i, \mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{S}\mathcal{C}_{\mathsf{pr}}^{1/2}e_i \Big\rangle_{\mathscr{H}} = \Big\langle \mathcal{C}_{\mathsf{pr}}^{1/2}e_i, \mathcal{S}\mathcal{C}_{\mathsf{pr}}^{1/2}e_i \Big\rangle_{\mathscr{H}} \\
&= \sum_j (1+\lambda_j)^{-1} \Big\langle e_j, \mathcal{C}_{\mathsf{pr}}^{1/2}e_i \Big\rangle_{\mathscr{H}}^2 \leq \sum_j \Big\langle e_j, \mathcal{C}_{\mathsf{pr}}^{1/2}e_i \Big\rangle_{\mathscr{H}}^2 = \Big\| \mathcal{C}_{\mathsf{pr}}^{1/2}e_i \Big\|_{\mathscr{H}}^2 = \langle e_i, \mathcal{C}_{\mathsf{pr}}e_i \rangle_{\mathscr{H}},
\end{aligned}$$

where the penultimate equality follows from Parseval's identity. □

Also,

$$\mathrm{tr}(\mathcal{C}_{\mathsf{post}}) = \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}) - \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}^{1/2}(I - \mathcal{S})\mathcal{C}_{\mathsf{pr}}^{1/2}) = \sum_{j=1}^{\infty} (1 - \alpha_j) \langle e_j, \mathcal{C}_{\mathsf{pr}}e_j \rangle_{\mathscr{H}},$$

where $\alpha_j = \lambda_j/(1 + \lambda_j)$. Thus, for eigenvalues λ_j that are large, we have $\alpha_j \approx 1$ which suggests that significant uncertainty reduction occurs in such directions. It is well known that for large classes of ill-posed Bayesian inverse problems, the eigenvalues λ_i of $\tilde{\mathcal{H}}_{\mathsf{m}}$ decay rapidly to zero, with a relatively small number of dominant eigenvalues

indicating the *data-informed* directions in the parameter space. This allows "focusing" the inference to low-dimensional subspaces of the parameter space $\mathscr{H}$. Such ideas have been used to develop efficient numerical algorithms for solution of infinite-dimensional Bayesian inverse problems in works such as [4, 11] and for algorithms for computing A-optimal experimental designs for infinite-dimensional Bayesian linear inverse problems in [1].

4. KL divergence from posterior to prior and expected information gain

Let us first motivate the discussion by recalling the form of the KL divergence from the posterior to prior in the finite-dimensional case. We use boldface letters for the finite-dimensional versions of the operators appearing in the Bayesian inverse problem. To indicate that we work in $\mathbb{R}^n$, we denote by $\mu_{\mathsf{pr,n}}$ and $\mu^{\boldsymbol{y}}_{\mathsf{post,n}}$ the prior and posterior measures in the n-dimensional case. The following expression for $D_{\mathsf{kl}}\left(\mu^{\boldsymbol{y}}_{\mathsf{post,n}} \| \mu_{\mathsf{pr,n}}\right)$ is well known:

$$D_{\mathsf{kl}}\left(\mu^{\boldsymbol{y}}_{\mathsf{post,n}} \| \mu_{\mathsf{pr,n}}\right) = \frac{1}{2}\Big[-\log\left(\frac{\det \mathbf{C}_{\mathsf{post}}}{\det \mathbf{C}_{\mathsf{pr}}}\right) - n + \operatorname{tr}(\mathbf{C}_{\mathsf{pr}}^{-1}\mathbf{C}_{\mathsf{post}}) + \left\langle \mathbf{C}_{\mathsf{pr}}^{-1}\boldsymbol{u}^{\boldsymbol{y}}_{\mathsf{post}}, \boldsymbol{u}^{\boldsymbol{y}}_{\mathsf{post}}\right\rangle_{\mathbb{R}^n}\Big]. \quad (11)$$

Note that the above expression is not meaningful in the infinite-dimensional case. For one thing, n appears explicitly in the expression. Moreover, in the infinite-dimensional case, $\mathcal{C}_{\mathsf{pr}}$ is a trace-class operator whose eigenvalues accumulate at zero, so dividing by the determinant of the prior covariance is problematic as $n \to \infty$. Finally, in the infinite-dimensional case, $\mathcal{C}_{\mathsf{pr}}^{-1}$ is the inverse of a compact operator and hence is unbounded; therefore, the trace term, which involves the inverse of the prior covariance, needs clarification. However, if we reformulate the above expression, we obtain an expression that has meaning in the infinite-dimensional case.

A straightforward calculation shows that the first term on the right in (11) may be simplified:

$$\begin{aligned} -\log\left(\frac{\det \mathbf{C}_{\mathsf{post}}}{\det \mathbf{C}_{\mathsf{pr}}}\right) \;&=\; \log\left(\frac{\det \mathbf{C}_{\mathsf{pr}}}{\det \mathbf{C}_{\mathsf{post}}}\right) = \log\det\left(\mathbf{C}_{\mathsf{pr}}\mathbf{C}_{\mathsf{post}}^{-1}\right) \\ &= \log\det\left(\mathbf{C}_{\mathsf{pr}}^{1/2}(\mathbf{H}_{\mathsf{m}} + \mathbf{C}_{\mathsf{pr}}^{-1})\mathbf{C}_{\mathsf{pr}}^{1/2}\right) \\ &= \log\det(\tilde{\mathbf{H}}_{\mathsf{m}} + \mathbf{I}). \end{aligned} \quad (12)$$

Recall that, in general, if $\mathbf{A}$ is Hermitian, then there exists a unitary matrix $\mathbf{U}$ such that

$$\mathbf{D} = [\lambda_i \delta_{ij}] = \mathbf{U}^*\mathbf{A}\mathbf{U}$$

is diagonal. In this case, the diagonal elements are the eigenvalues of $\mathbf{A}$, and

$$\det(\mathbf{I} + \mathbf{A}) = \det(\mathbf{U})\det(\mathbf{I} + \mathbf{A})\det(\mathbf{U}^*) = \prod_{i=1}^{n}(1 + \lambda_i).$$

In the infinite-dimensional setting, given a trace-class operator $\mathcal{A} \in \mathscr{L}_1^{\mathsf{sym}+}(\mathscr{H})$,

$$\lim_{n\to\infty} \log\left(\prod_{i=1}^{n}(1 + \lambda_i(\mathcal{A}))\right) = \lim_{n\to\infty}\sum_{i=1}^{n}\log(1 + \lambda_i(\mathcal{A})) \le \lim_{n\to\infty}\sum_{i=1}^{n}\lambda_i(\mathcal{A}) < \infty,$$

so, motivated by the n-dimensional case, we may define the Fredholm determinant of $I + \mathcal{A}$ as

$$\det(I + \mathcal{A}) = \prod_{i=1}^{\infty}(1 + \lambda_i(\mathcal{A})),$$

where $\lambda_i(\mathcal{A})$ are the eigenvalues of $\mathcal{A}$ [15]. Hence, the final expression in equation (12) is meaningful in infinite dimensions. Next, we consider the term $-n + \mathrm{tr}(\mathbf{C}_{\mathsf{pr}}^{-1}\mathbf{C}_{\mathsf{post}})$:

$$\begin{aligned}-n + \mathrm{tr}(\mathbf{C}_{\mathsf{pr}}^{-1}\mathbf{C}_{\mathsf{post}}) &= -\mathrm{tr}(\mathbf{I}) + \mathrm{tr}(\mathbf{C}_{\mathsf{pr}}^{-1}\mathbf{C}_{\mathsf{post}}) \\ &= \mathrm{tr}(\mathbf{C}_{\mathsf{pr}}^{-1}\mathbf{C}_{\mathsf{post}} - \mathbf{I}) = \mathrm{tr}\big((\mathbf{C}_{\mathsf{pr}}^{-1} - \mathbf{C}_{\mathsf{post}}^{-1})\mathbf{C}_{\mathsf{post}}\big) = -\mathrm{tr}(\mathbf{H}_{\mathsf{m}}\mathbf{C}_{\mathsf{post}}),\end{aligned}$$

where in the last step we used the fact that $\mathbf{C}_{\mathsf{post}}^{-1} = \mathbf{H}_{\mathsf{m}} + \mathbf{C}_{\mathsf{pr}}^{-1}$. Notice that the argument of the trace in the final expression is in fact a trace-class operator in the infinite-dimensional case and has a well-defined trace. Combining (12) and (13) and defining the inner-product $\langle \boldsymbol{x}, \boldsymbol{y}\rangle_{\mathbf{C}_{\mathsf{pr}}^{-1}} = \left\langle \mathbf{C}_{\mathsf{pr}}^{-1/2}\boldsymbol{x}, \mathbf{C}_{\mathsf{pr}}^{-1/2}\boldsymbol{y}\right\rangle_{\mathbb{R}^n}$ for $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$, we rewrite (11),

$$D_{\mathsf{kl}}\left(\mu_{\mathsf{post,n}}^{\boldsymbol{y}} \| \mu_{\mathsf{pr,n}}\right) = \frac{1}{2}\Big[\log\det(\tilde{\mathbf{H}}_{\mathsf{m}} + \mathbf{I}) - \mathrm{tr}(\mathbf{H}_{\mathsf{m}}\mathbf{C}_{\mathsf{post}}) + \left\langle \boldsymbol{u}_{\mathsf{post}}^{\boldsymbol{y}}, \boldsymbol{u}_{\mathsf{post}}^{\boldsymbol{y}}\right\rangle_{\mathbf{C}_{\mathsf{pr}}^{-1}}\Big]. \quad (13)$$

In Section 4.1 we derive, rigorously, alternate forms of the expression for the KL divergence from posterior to prior in the infinite-dimensional Hilbert space setting; as we shall see shortly, one of those forms is a direct extension of (13) to the infinite-dimensional case. The reason for introducing the weighted inner-product $\langle\cdot,\cdot\rangle_{\mathbf{C}_{\mathsf{pr}}^{-1}}$ will also become clear in the discussion that follows.

4.1. **The KL-divergence from posterior to prior.** The following result which is a consequence of Proposition 1.2.8 in [9] will be needed in what follows.

Proposition 2. *Let $\mathcal{A} \in \mathscr{L}(\mathscr{H})$ be a positive self-adjoint operator, $\mu = \mathcal{N}(0, \mathcal{Q})$ a non-degenerate Gaussian measure on $\mathscr{H}$, and $b \in \mathscr{H}$. Then,*

$$\int_{\mathscr{H}} \exp\left\{-\frac{1}{2}\left\langle \mathcal{A}x, x\right\rangle_{\mathscr{H}} + \left\langle b, x\right\rangle_{\mathscr{H}}\right\}\mu(dx) = \det(I+\tilde{\mathcal{A}})^{-1/2}\exp\left\{\frac{1}{2}\left\|(I+\tilde{\mathcal{A}})^{-1/2}\mathcal{Q}^{1/2}b\right\|_{\mathscr{H}}^2\right\},$$

where $\tilde{\mathcal{A}} = \mathcal{Q}^{1/2}\mathcal{A}\mathcal{Q}^{1/2}$.

In the following technical lemma, we calculate the expression for $\mathcal{Z}_0$, introduced in equation (6).

Lemma 2. *Let $\Phi(u; \boldsymbol{y}) = \frac{1}{2}(\mathcal{G}u - \boldsymbol{y})^T\boldsymbol{\Gamma}_{\mathit{noise}}^{-1}(\mathcal{G}u - \boldsymbol{y})$, as defined by equation (4). Then,*

$$\mathcal{Z}_0(\boldsymbol{y}) := \int_{\mathscr{H}} \exp\{-\Phi(u;\boldsymbol{y})\}\,\mu_{\mathit{pr}}(du) = \exp\left\{-\frac{1}{2}|\boldsymbol{y}|^2\right\}\det(I+\tilde{\mathcal{H}}_{\mathit{m}})^{-1/2}\exp\left\{\frac{1}{2}\left\langle \mathcal{C}_{\mathit{post}}b, b\right\rangle_{\mathscr{H}}\right\},$$

where $b = \mathcal{G}^*\boldsymbol{y}$ and $\mathcal{C}_{\mathit{post}} = (\mathcal{G}^*\mathcal{G} + \mathcal{C}_{\mathit{pr}}^{-1})^{-1}$, as in equation (7).

Proof. First note that (recall that we have assumed $\boldsymbol{\Gamma}_{\mathsf{noise}} = \mathbf{I}$)

$$\begin{aligned}\Phi(u;\boldsymbol{y}) \;&=\; \frac{1}{2}(\mathcal{G}u - \boldsymbol{y})^T(\mathcal{G}u - \boldsymbol{y}) = \frac{1}{2}\left\langle \mathcal{G}u, \mathcal{G}u\right\rangle_{\mathbb{R}^q} - \left\langle \mathcal{G}u, \boldsymbol{y}\right\rangle_{\mathbb{R}^q} + \frac{1}{2}\left\langle \boldsymbol{y}, \boldsymbol{y}\right\rangle_{\mathbb{R}^q} \\ &= \frac{1}{2}\left\langle \mathcal{H}_{\mathsf{m}}u, u\right\rangle_{\mathscr{H}} - \left\langle \mathcal{G}^*\boldsymbol{y}, u\right\rangle_{\mathscr{H}} + \frac{1}{2}|\boldsymbol{y}|^2. \qquad (14)\end{aligned}$$

Therefore,

$$\int_{\mathscr{H}} \exp\{-\Phi(u;\boldsymbol{y})\}\,\mu_{\mathsf{pr}}(du) = \exp\left\{-\frac{1}{2}|\boldsymbol{y}|^2\right\}\int_{\mathscr{H}}\exp\left\{-\frac{1}{2}\left\langle \mathcal{H}_{\mathsf{m}}u, u\right\rangle_{\mathscr{H}} + \left\langle b, u\right\rangle_{\mathscr{H}}\right\}\mu_{\mathsf{pr}}(du),$$

where $b = \mathcal{G}^* \boldsymbol{y}$. By Proposition 2 we have,

$$\begin{aligned}\int_{\mathscr{H}} \exp\left\{-\frac{1}{2}\left\langle \mathcal{H}_{\mathsf{m}} u, u\right\rangle_{\mathscr{H}} + \left\langle b, u\right\rangle_{\mathscr{H}}\right\} \mu_{\mathsf{pr}}(du) \\ = \det(I + \mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}^{1/2})^{-1/2} \exp\left\{\frac{1}{2}\left\|(I + \mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}^{1/2})^{-1/2}\mathcal{C}_{\mathsf{pr}}^{1/2} b\right\|_{\mathscr{H}}^2\right\} \\ = \det(I + \tilde{\mathcal{H}}_{\mathsf{m}})^{-1/2} \exp\left\{\frac{1}{2}\left\|(I + \tilde{\mathcal{H}}_{\mathsf{m}})^{-1/2}\mathcal{C}_{\mathsf{pr}}^{1/2} b\right\|_{\mathscr{H}}^2\right\}.\end{aligned}$$

The assertion of the lemma now follows, since $\mathcal{C}_{\mathsf{post}} = \mathcal{C}_{\mathsf{pr}}^{1/2}(I + \tilde{\mathcal{H}}_{\mathsf{m}})^{-1}\mathcal{C}_{\mathsf{pr}}^{1/2}$. □

The following result provides the expression for the KL divergence from posterior to prior:

Proposition 3. *Let μ_{pr} be a centered Gaussian measure on $\mathscr{H}$, and $\mu_{post}^{\boldsymbol{y}} = \mathcal{N}(u_{post}^{\boldsymbol{y}}, \mathcal{C}_{post})$ be the posterior measure for a Bayesian linear inverse problem with additive Gaussian noise model as described in Section 3. Then,*

$$D_{kl}\left(\mu_{post}^{\boldsymbol{y}} \| \mu_{pr}\right) = \frac{1}{2}\left[\log\det(I + \tilde{\mathcal{H}}_m) - \mathrm{tr}(\mathcal{H}_m \mathcal{C}_{post}) - \left\langle u_{post}^{\boldsymbol{y}}, \mathcal{G}^*(\mathcal{G} u_{post}^{\boldsymbol{y}} - \boldsymbol{y})\right\rangle_{\mathscr{H}}\right]. \tag{15}$$

Proof. Consider (6), and note that

$$\begin{aligned} D_{\mathsf{kl}}\left(\mu_{\mathsf{post}}^{\boldsymbol{y}} \| \mu_{\mathsf{pr}}\right) &= \int_{\mathscr{H}} \log\left\{\frac{d\mu_{\mathsf{post}}^{\boldsymbol{y}}}{d\mu_{\mathsf{pr}}}\right\} \mu_{\mathsf{post}}^{\boldsymbol{y}}(du) \\ &= -\log \mathcal{Z}_0(\boldsymbol{y}) - \int_{\mathscr{H}} \Phi(u; \boldsymbol{y})\, \mu_{\mathsf{post}}^{\boldsymbol{y}}(du). \end{aligned} \tag{16}$$

Using (14) to expand $\Phi(u; \boldsymbol{y})$, the integral on the right becomes

$$\int_{\mathscr{H}} \Phi(u; \boldsymbol{y})\, \mu_{\mathsf{post}}^{\boldsymbol{y}}(du) = \frac{1}{2}\int_{\mathscr{H}} \left\langle \mathcal{H}_{\mathsf{m}} u, u\right\rangle_{\mathscr{H}} \mu_{\mathsf{post}}^{\boldsymbol{y}}(du) - \int_{\mathscr{H}} \left\langle \mathcal{G}^* \boldsymbol{y}, u\right\rangle_{\mathscr{H}} \mu_{\mathsf{post}}^{\boldsymbol{y}}(du) + \frac{1}{2}|\boldsymbol{y}|^2.$$

The second integral evaluates to $\left\langle \mathcal{G}^* \boldsymbol{y}, u_{\mathsf{post}}^{\boldsymbol{y}}\right\rangle_{\mathscr{H}}$, by the definition of the mean of the measure, and the first integral is evaluated via the formula for the integral of a quadratic form:

$$\int_{\mathscr{H}} \left\langle \mathcal{H}_{\mathsf{m}} u, u\right\rangle_{\mathscr{H}} \mu_{\mathsf{post}}^{\boldsymbol{y}}(du) = \mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}) + \left\langle u_{\mathsf{post}}^{\boldsymbol{y}}, \mathcal{H}_{\mathsf{m}} u_{\mathsf{post}}^{\boldsymbol{y}}\right\rangle_{\mathscr{H}}.$$

Using the expression for $\mathcal{Z}_0$ from Lemma 2,

$$\begin{aligned} -\log \mathcal{Z}_0(\boldsymbol{y}) &= \frac{1}{2}|\boldsymbol{y}|^2 - \log\det(I + \tilde{\mathcal{H}}_{\mathsf{m}})^{-1/2} - \frac{1}{2}\left\langle \mathcal{C}_{\mathsf{post}}\mathcal{G}^*\boldsymbol{y}, \mathcal{G}^*\boldsymbol{y}\right\rangle_{\mathscr{H}} \\ &= \frac{1}{2}|\boldsymbol{y}|^2 + \frac{1}{2}\log\det(I + \tilde{\mathcal{H}}_{\mathsf{m}}) - \frac{1}{2}\left\langle u_{\mathsf{post}}^{\boldsymbol{y}}, \mathcal{G}^*\boldsymbol{y}\right\rangle_{\mathscr{H}}, \end{aligned}$$

where we have also used the definition of $u_{\mathsf{post}}^{\boldsymbol{y}}$. Substituting into equation (16), we obtain

$$\begin{aligned} D_{\mathsf{kl}}\left(\mu_{\mathsf{post}}^{\boldsymbol{y}} \| \mu_{\mathsf{pr}}\right) =& \frac{1}{2}\log\det(I + \tilde{\mathcal{H}}_{\mathsf{m}}) - \frac{1}{2}\left\langle u_{\mathsf{post}}^{\boldsymbol{y}}, \mathcal{G}^*\boldsymbol{y}\right\rangle_{\mathscr{H}} \\ &- \frac{1}{2}\mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}) - \frac{1}{2}\left\langle u_{\mathsf{post}}^{\boldsymbol{y}}, \mathcal{H}_{\mathsf{m}} u_{\mathsf{post}}^{\boldsymbol{y}}\right\rangle_{\mathscr{H}} + \left\langle \mathcal{G}^*\boldsymbol{y}, u_{\mathsf{post}}^{\boldsymbol{y}}\right\rangle_{\mathscr{H}}, \end{aligned}$$

which, after some algebraic manipulation and recalling that $\mathcal{H}_{\mathsf{m}} = \mathcal{G}^*\mathcal{G}$, yields the assertion of the proposition. □

Let us note the following interpretation for the last term appearing in $D_{\mathsf{kl}}\left(\mu_{\mathsf{post}}^{\boldsymbol{y}}\|\mu_{\mathsf{pr}}\right)$ given in (15). Consider the function $\Phi(u) = \frac{1}{2}(\mathcal{G}u - \boldsymbol{y})^T(\mathcal{G}u - \boldsymbol{y})$, which is the familiar misfit term in the deterministic interpretation of the corresponding linear inverse problem. (For notational simplicity we have suppressed the dependence of Φ on the data vector $\boldsymbol{y}$.) Note that the variational derivative of Φ at a point $u \in \mathscr{H}$ in direction $h \in \mathscr{H}$ is given by,

$$\Phi'(u)h = \frac{d}{d\varepsilon}\Big|_{\varepsilon=0} \Phi(u + \varepsilon h) = \langle \mathcal{G}u - \boldsymbol{y}, \mathcal{G}h\rangle_{\mathbb{R}^q} = \langle \mathcal{G}^*(\mathcal{G}u - \boldsymbol{y}), h\rangle_{\mathscr{H}}\,.$$

Next, recall that the mean of the posterior, $u_{\mathsf{post}}^{\boldsymbol{y}}$, of the present Bayesian linear inverse problem coincides with the MAP estimator for the inference parameter u and is the global minimizer of the following regularized cost functional [16, 10]

$$\mathcal{J}(u) = \Phi(u) + \frac{1}{2}\,\langle u, u\rangle_{\mathcal{C}_{\mathsf{pr}}^{-1}}$$

with minimization done over the space, $\mathscr{H}_{\mu_{\mathsf{pr}}} = \mathsf{range}(\mathcal{C}_{\mathsf{pr}}^{1/2}) \subset \mathscr{H}$.[2] The inner-product in the regularization term is given by $\langle\cdot,\cdot\rangle_{\mathcal{C}_{\mathsf{pr}}^{-1}} = \left\langle \mathcal{C}_{\mathsf{pr}}^{-1/2}x, \mathcal{C}_{\mathsf{pr}}^{-1/2}y\right\rangle_{\mathscr{H}}$ for $x, y \in \mathscr{H}_{\mu_{\mathsf{pr}}}$. We have, by the first order optimality conditions $\mathcal{J}'(u_{\mathsf{post}}^{\boldsymbol{y}})h = 0$ for every $h \in \mathscr{H}_{\mu_{\mathsf{pr}}}$, that is,

$$\langle \mathcal{G}^*(\mathcal{G}u - \boldsymbol{y}), h\rangle_{\mathscr{H}} + \left\langle u_{\mathsf{post}}^{\boldsymbol{y}}, h\right\rangle_{\mathcal{C}_{\mathsf{pr}}^{-1}} = 0, \qquad \text{for all } h \in \mathscr{H}_{\mu_{\mathsf{pr}}}.$$

Thus, in particular, $-\left\langle \mathcal{G}^*(\mathcal{G}u - \boldsymbol{y}), u_{\mathsf{post}}^{\boldsymbol{y}}\right\rangle_{\mathscr{H}} = \left\langle u_{\mathsf{post}}^{\boldsymbol{y}}, u_{\mathsf{post}}^{\boldsymbol{y}}\right\rangle_{\mathcal{C}_{\mathsf{pr}}^{-1}}$. This leads to the following alternate form of expression (15):

$$D_{\mathsf{kl}}\left(\mu_{\mathsf{post}}^{\boldsymbol{y}}\|\mu_{\mathsf{pr}}\right) = \frac{1}{2}\left[\log\det(I + \tilde{\mathcal{H}}_{\mathsf{m}}) - \mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}) + \left\langle u_{\mathsf{post}}^{\boldsymbol{y}}, u_{\mathsf{post}}^{\boldsymbol{y}}\right\rangle_{\mathcal{C}_{\mathsf{pr}}^{-1}}\right]. \tag{17}$$

Note that this expression for the KL divergence $D_{\mathsf{kl}}\left(\mu_{\mathsf{post}}^{\boldsymbol{y}}\|\mu_{\mathsf{pr}}\right)$ is the direct extension of the corresponding expression in the case of $\mathscr{H} = \mathbb{R}^n$ as given in (13) to infinite dimensions.

Remark 1. A straightforward modification of the arguments leading to equation (15), for the case of a prior $\mu_{\mathsf{pr}} = \mathcal{N}(u_{\mathsf{pr}}, \mathcal{C}_{\mathsf{pr}})$, leads to

$$D_{\mathsf{kl}}\left(\mu_{\mathsf{post}}^{\boldsymbol{y}}\|\mu_{\mathsf{pr}}\right) = \frac{1}{2}\left[\log\det(I + \tilde{\mathcal{H}}_{\mathsf{m}}) - \mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}) - \left\langle u_{\mathsf{post}}^{\boldsymbol{y}} - u_{\mathsf{pr}}, \mathcal{G}^*(\mathcal{G}u_{\mathsf{post}}^{\boldsymbol{y}} - \boldsymbol{y})\right\rangle_{\mathscr{H}}\right].$$

Moreover, in view of the argument leading to (17), we have:

$$D_{\mathsf{kl}}\left(\mu_{\mathsf{post}}^{\boldsymbol{y}}\|\mu_{\mathsf{pr}}\right) = \frac{1}{2}\left[\log\det(I + \tilde{\mathcal{H}}_{\mathsf{m}}) - \mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}) + \left\langle u_{\mathsf{post}}^{\boldsymbol{y}} - u_{\mathsf{pr}}, u_{\mathsf{post}}^{\boldsymbol{y}} - u_{\mathsf{pr}}\right\rangle_{\mathcal{C}_{\mathsf{pr}}^{-1}}\right].$$

4.2. **Expected information gain.** Here we derive the expression for the expected information gain. We first prove the following technical lemma which is needed in the proof of the main result in this section.

Lemma 3. *The following identities hold.*

(1) $\mathsf{E}_{\mu_{pr}}\left\{\mathsf{E}_{\boldsymbol{y}|u}\left\{\left\langle u_{post}^{\boldsymbol{y}}, \mathcal{G}^*\boldsymbol{y}\right\rangle_{\mathscr{H}}\right\}\right\} = \mathrm{tr}(\tilde{\mathcal{H}}_m)$

(2) $\mathsf{E}_{\mu_{pr}}\left\{\mathsf{E}_{\boldsymbol{y}|u}\left\{\left\langle u_{post}^{\boldsymbol{y}}, \mathcal{H}_m u_{post}^{\boldsymbol{y}}\right\rangle_{\mathscr{H}}\right\}\right\} = \mathrm{tr}(\mathcal{S}\tilde{\mathcal{H}}_m^2)$,

where $\tilde{\mathcal{H}}_m$ *and* $\mathcal{S}$ *be as in* (9) *and* (10) *respectively.*

[2]Given a Gaussian measure $\mu = \mathcal{N}(m, \mathcal{C})$ on a Hilbert space H, the space $\mathsf{range}(\mathcal{C}^{1/2})$ is called the Cameron-Martin space corresponding to the measure μ. It is a known result (see e.g. [8]) that if the Hilbert space H is infinite-dimensional, $\mu\big(\mathsf{range}(\mathcal{C}^{1/2})\big) = 0$.

Proof. We present the proof of the first statement; the second one follows from a similar argument. Let us begin from the inner expectation. Note that, by the definition of $u^{\boldsymbol{y}}_{\mathsf{post}}$ we have,

$$\left\langle u^{\boldsymbol{y}}_{\mathsf{post}}, \mathcal{G}^*\boldsymbol{y}\right\rangle_{\mathscr{H}} = \langle \mathcal{C}_{\mathsf{post}}\mathcal{G}^*\boldsymbol{y}, \mathcal{G}^*\boldsymbol{y}\rangle_{\mathscr{H}} = \langle \boldsymbol{y}, \mathcal{G}\mathcal{C}_{\mathsf{post}}\mathcal{G}^*\boldsymbol{y}\rangle_{\mathbb{R}^q},$$

For clarity let us denote $L = \mathcal{G}\mathcal{C}_{\mathsf{post}}\mathcal{G}^*$. Recall that $\boldsymbol{y}|u$ is distributed according to $\mathcal{N}(\mathcal{G}u, \boldsymbol{\Gamma}_{\mathsf{noise}})$, and that we assumed $\boldsymbol{\Gamma}_{\mathsf{noise}} = \mathbf{I}$. Using the formula for the expectation of a quadratic form (on $\mathscr{Y} = \mathbb{R}^q$), Lemma 1, we have

$$\mathsf{E}_{\boldsymbol{y}|u}\left\{\left\langle u^{\boldsymbol{y}}_{\mathsf{post}}, \mathcal{G}^*\boldsymbol{y}\right\rangle_{\mathscr{H}}\right\} = \mathsf{E}_{\boldsymbol{y}|u}\left\{\langle \boldsymbol{y}, L\boldsymbol{y}\rangle_{\mathbb{R}^q}\right\} = \mathrm{tr}(L) + \langle \mathcal{G}u, L\mathcal{G}u\rangle_{\mathbb{R}^q} = \mathrm{tr}(L) + \langle u, \mathcal{G}^*L\mathcal{G}u\rangle_{\mathscr{H}}.$$

By the comment at the end of Section 2.1 and recalling that $\mathcal{C}_{\mathsf{post}} = \mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{S}\mathcal{C}_{\mathsf{pr}}^{1/2}$, we have

$$\begin{aligned}\mathrm{tr}(L) \;&=\; \mathrm{tr}(\mathcal{G}\mathcal{C}_{\mathsf{post}}\mathcal{G}^*) = \mathrm{tr}(\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}}) = \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{S}\mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{H}_{\mathsf{m}}) \\ &= \mathrm{tr}(\mathcal{S}\mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}^{1/2}) = \mathrm{tr}(\mathcal{S}\tilde{\mathcal{H}}_{\mathsf{m}}). \end{aligned} \tag{18}$$

Therefore,

$$\mathsf{E}_{\boldsymbol{y}|u}\left\{\left\langle u^{\boldsymbol{y}}_{\mathsf{post}}, \mathcal{G}^*\boldsymbol{y}\right\rangle_{\mathscr{H}}\right\} = \mathrm{tr}(\mathcal{S}\tilde{\mathcal{H}}_{\mathsf{m}}) + \langle u, \mathcal{G}^*L\mathcal{G}u\rangle_{\mathscr{H}}. \tag{19}$$

Next, to compute the outer expectation we proceed as follows (keep in mind that $\mu_{\mathsf{pr}} = \mathcal{N}(0, \mathcal{C}_{\mathsf{pr}})$). By Lemma 1,

$$\mathsf{E}_{\mu_{\mathsf{pr}}}\left\{\langle u, \mathcal{G}^*L\mathcal{G}u\rangle_{\mathscr{H}}\right\} = \int_{\mathscr{H}} \langle u, \mathcal{G}^*L\mathcal{G}u\rangle_{\mathscr{H}}\; \mu_{\mathsf{pr}}(du) = \mathrm{tr}(\mathcal{G}^*L\mathcal{G}\mathcal{C}_{\mathsf{pr}});$$

and

$$\begin{aligned}\mathrm{tr}(\mathcal{G}^*L\mathcal{G}\mathcal{C}_{\mathsf{pr}}) \;&=\; \mathrm{tr}(\mathcal{G}^*\mathcal{G}\mathcal{C}_{\mathsf{post}}\mathcal{G}^*\mathcal{G}\mathcal{C}_{\mathsf{pr}}) = \mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}) \\ &= \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}^{1/2}) = \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{S}\mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}^{1/2}) = \mathrm{tr}(\tilde{\mathcal{H}}_{\mathsf{m}}\mathcal{S}\tilde{\mathcal{H}}_{\mathsf{m}}) = \mathrm{tr}(\mathcal{S}\tilde{\mathcal{H}}_{\mathsf{m}}^2).\end{aligned}$$

Thus, combining equations (18), (19), and (20) gives

$$\begin{aligned}\mathsf{E}_{\mu_{\mathsf{pr}}}\left\{\mathsf{E}_{\boldsymbol{y}|u}\left\{\left\langle u^{\boldsymbol{y}}_{\mathsf{post}}, \mathcal{G}^*\boldsymbol{y}\right\rangle_{\mathscr{H}}\right\}\right\} \;&=\; \mathrm{tr}(\mathcal{S}\tilde{\mathcal{H}}_{\mathsf{m}}) + \mathrm{tr}(\mathcal{S}\tilde{\mathcal{H}}_{\mathsf{m}}^2) = \mathrm{tr}(\mathcal{S}\tilde{\mathcal{H}}_{\mathsf{m}}(I + \tilde{\mathcal{H}}_{\mathsf{m}})) \\ &= \mathrm{tr}(\tilde{\mathcal{H}}_{\mathsf{m}}(I + \tilde{\mathcal{H}}_{\mathsf{m}})\mathcal{S}) = \mathrm{tr}(\tilde{\mathcal{H}}_{\mathsf{m}}),\end{aligned}$$

which is the first statement of the lemma. □

The following theorem is the main result of this section.

Theorem 1. *Let μ_{pr} be a centered Gaussian prior measure on $\mathscr{H}$, and $\mu^{\boldsymbol{y}}_{post} = \mathcal{N}\big(u^{\boldsymbol{y}}_{post}, \mathcal{C}_{post}\big)$ be the posterior measure for a Bayesian linear inverse problem with additive Gaussian noise model as described in Section 3. Then,*

$$\mathsf{E}_{\mu_{pr}}\left\{\mathsf{E}_{\boldsymbol{y}|u}\left\{D_{kl}\left(\mu^{\boldsymbol{y}}_{post}\|\mu_{pr}\right)\right\}\right\} = \frac{1}{2}\log\det(I + \tilde{\mathcal{H}}_m).$$

Proof. By (15) we have,

$$\begin{aligned}\mathsf{E}_{\mu_{\mathsf{pr}}}\left\{\mathsf{E}_{\boldsymbol{y}|u}\left\{D_{\mathsf{kl}}\left(\mu^{\boldsymbol{y}}_{\mathsf{post}}\|\mu_{\mathsf{pr}}\right)\right\}\right\} \;&=\; \frac{1}{2}\log\det(I + \tilde{\mathcal{H}}_{\mathsf{m}}) \\ &\quad -\frac{1}{2}\mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}) - \frac{1}{2}\mathsf{E}_{\mu_{\mathsf{pr}}}\left\{\mathsf{E}_{\boldsymbol{y}|u}\left\{\left\langle u^{\boldsymbol{y}}_{\mathsf{post}}, \mathcal{G}^*(\mathcal{G}u^{\boldsymbol{y}}_{\mathsf{post}} - \boldsymbol{y})\right\rangle_{\mathscr{H}}\right\}\right\}. \end{aligned} \tag{20}$$

Using the previous lemma we can proceed as follows,

$$\begin{aligned}&\mathsf{E}_{\mu_{\mathsf{pr}}}\left\{\mathsf{E}_{\boldsymbol{y}|u}\left\{\left\langle u^{\boldsymbol{y}}_{\mathsf{post}}, \mathcal{G}^*(\mathcal{G}u^{\boldsymbol{y}}_{\mathsf{post}} - \boldsymbol{y})\right\rangle_{\mathscr{H}}\right\}\right\} \\ &\qquad = \mathsf{E}_{\mu_{\mathsf{pr}}}\left\{\mathsf{E}_{\boldsymbol{y}|u}\left\{\left\langle u^{\boldsymbol{y}}_{\mathsf{post}}, \mathcal{H}_{\mathsf{m}}u^{\boldsymbol{y}}_{\mathsf{post}}\right\rangle_{\mathscr{H}}\right\}\right\} - \mathsf{E}_{\mu_{\mathsf{pr}}}\left\{\mathsf{E}_{\boldsymbol{y}|u}\left\{\left\langle u^{\boldsymbol{y}}_{\mathsf{post}}, \mathcal{G}^*\boldsymbol{y}\right\rangle_{\mathscr{H}}\right\}\right\} \\ &\qquad = \mathrm{tr}(\mathcal{S}\tilde{\mathcal{H}}_{\mathsf{m}}^2) - \mathrm{tr}(\tilde{\mathcal{H}}_{\mathsf{m}}) \\ &\qquad = \mathrm{tr}(\mathcal{S}(\tilde{\mathcal{H}}_{\mathsf{m}} - \mathcal{S}^{-1})\tilde{\mathcal{H}}_{\mathsf{m}}) = -\mathrm{tr}(\mathcal{S}\tilde{\mathcal{H}}_{\mathsf{m}}).\end{aligned}$$

Thus, since $\mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}) = \mathrm{tr}(\tilde{\mathcal{H}}_{\mathsf{m}}\mathcal{S}) = \mathrm{tr}(\mathcal{S}\tilde{\mathcal{H}}_{\mathsf{m}})$, the expression for the expected information gain in (20) simplifies to $\mathsf{E}_{\mu_{\mathsf{pr}}}\left\{\mathsf{E}_{\boldsymbol{y}|u}\left\{D_{\mathsf{kl}}\left(\mu_{\mathsf{post}}^{\boldsymbol{y}} \| \mu_{\mathsf{pr}}\right)\right\}\right\} = \frac{1}{2}\log\det(I + \tilde{\mathcal{H}}_{\mathsf{m}})$. □

The result above provides the infinite-dimensional analogue of Bayesian D-optimality. As mentioned in the introduction, an experimental design $\boldsymbol{\xi}$ enters the Bayesian inverse problem through the data likelihood. This dependence to $\boldsymbol{\xi}$, in the present Gaussian linear case, is manifested through a $\boldsymbol{\xi}$ dependent misfit Hessian, $\mathcal{H}_{\mathsf{m}} = \mathcal{H}_{\mathsf{m}}(\boldsymbol{\xi})$. Consequently, the D-optimal design problem in the infinite-dimensional Hilbert space setting is given by,

$$\underset{\boldsymbol{\xi}\in\Xi}{\text{maximize}} \;\; \log\det(I + \tilde{\mathcal{H}}_{\mathsf{m}}(\boldsymbol{\xi})),$$

where Ξ is the design space which needs to be specified in a given experimental design problem.

Remark 2. As mentioned earlier, in a large class of Bayesian inverse problems, $\tilde{\mathcal{H}}_{\mathsf{m}}$ admits a low-rank approximation,

$$\tilde{\mathcal{H}}_{\mathsf{m}} v \approx \sum_{i=1}^{r} \lambda_i \left\langle e_i, v\right\rangle_{\mathscr{H}} e_i, \quad v \in \mathscr{H},$$

where r is the numerical rank of $\tilde{\mathcal{H}}_{\mathsf{m}}$ and $\{\lambda_i\}_{i=1}^r$ are the dominant eigenvalues of $\tilde{\mathcal{H}}_{\mathsf{m}}$ with respective eigenvectors $\{e_i\}_{i=1}^r$. Thus, one can use the following approximation

$$\log\det(I + \tilde{\mathcal{H}}_{\mathsf{m}}) \approx \sum_{i=1}^{r} \log(1+\lambda_i),$$

which enables an efficient means of approximating the expected information gain.

5. Expected mean square error of the MAP estimator and Bayesian A-optimality

In this section, we consider another well known optimal experimental design criterion, Bayesian A-optimality, which aims to minimize the trace of the posterior covariance operator. It is well known in the statistics literature that for inference problems with a finite-dimensional parameter, this is equivalent to minimizing the expected mean square error of the mean posterior which, in the case of a Bayesian linear inverse problem, coincides with the MAP estimator. In this section, we extend this result to the infinite-dimensional Hilbert space setting.

The MSE of the MAP estimator $u_{\mathsf{post}}^{\boldsymbol{y}}$ is

$$\mathrm{MSE}(u_{\mathsf{post}}^{\boldsymbol{y}}; u) = \mathsf{E}_{\boldsymbol{y}|u}\left\{\left\|u - u_{\mathsf{post}}^{\boldsymbol{y}}\right\|_{\mathscr{H}}^2\right\}.$$

The MSE is also referred to as the *risk* of the estimator $u_{\mathsf{post}}^{\boldsymbol{y}}$, corresponding to a quadratic loss function. A straightforward calculation shows that

$$\mathrm{MSE}(u_{\mathsf{post}}^{\boldsymbol{y}}; u) = \left\|u - \mathsf{E}_{\boldsymbol{y}|u}\left\{u_{\mathsf{post}}^{\boldsymbol{y}}\right\}\right\|_{\mathscr{H}}^2 + \mathsf{E}_{\boldsymbol{y}|u}\left\{\left\|u_{\mathsf{post}}^{\boldsymbol{y}} - \mathsf{E}_{\boldsymbol{y}|u}\left\{u_{\mathsf{post}}^{\boldsymbol{y}}\right\}\right\|_{\mathscr{H}}^2\right\}, \quad (21)$$

Note that the first term in (21) quantifies the magnitude of estimation bias, and the second term describes the variability of the estimator around its mean. The following technical Lemma provides the expression for $\mathrm{MSE}(u_{\mathsf{post}}^{\boldsymbol{y}}; u)$ in the infinite-dimensional Hilbert space setting.

Lemma 4. *Let $u_{post}^{\boldsymbol{y}}$ be the MAP estimator for u as in* (7)*. Then,*

$$\mathrm{MSE}(u_{post}^{\boldsymbol{y}}; u) = \|(\mathcal{C}_{post}\mathcal{H}_m - I)u\|_{\mathscr{H}}^2 + \mathrm{tr}(\mathcal{C}_{post}^2\mathcal{H}_m).$$

Proof. Consider the expression for $\mathrm{MSE}(u_{\mathsf{post}}^{\boldsymbol{y}}; u)$ given in (21). For the first term in the sum, we have

$$u - \mathsf{E}_{\boldsymbol{y}|u}\left\{u_{\mathsf{post}}^{\boldsymbol{y}}\right\} = u - \mathsf{E}_{\boldsymbol{y}|u}\left\{\mathcal{C}_{\mathsf{post}}\mathcal{G}^*\boldsymbol{y}\right\} = u - \mathcal{C}_{\mathsf{post}}\mathcal{G}^*\mathcal{G}u = (I - \mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}})u.$$

Next, note that $\xi(\boldsymbol{y}) = u_{\mathsf{post}}^{\boldsymbol{y}} - \mathsf{E}_{\boldsymbol{y}|u}\left\{u_{\mathsf{post}}^{\boldsymbol{y}}\right\}$ has law $\mu = \mathcal{N}(0, \mathcal{Q})$ with $\mathcal{Q} = (\mathcal{C}_{\mathsf{post}}\mathcal{G}^*)(\mathcal{C}_{\mathsf{post}}\mathcal{G}^*)^* = \mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}$. Therefore,

$$\mathsf{E}_{\boldsymbol{y}|u}\left\{\left\|u_{\mathsf{post}}^{\boldsymbol{y}} - \mathsf{E}_{\boldsymbol{y}|u}\left\{u_{\mathsf{post}}^{\boldsymbol{y}}\right\}\right\|_{\mathscr{H}}^2\right\} = \int_{\mathscr{H}} \|\xi\|_{\mathscr{H}}^2 \; \mu(d\xi) = \mathrm{tr}(\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}) = \mathrm{tr}(\mathcal{C}_{\mathsf{post}}^2\mathcal{H}_{\mathsf{m}}).$$

□

Next, we consider the average over the prior measure of the MSE,

$$\mathsf{E}_{\mu_{\mathsf{pr}}}\left\{\mathrm{MSE}(u_{\mathsf{post}}^{\boldsymbol{y}}; u)\right\} = \int_{\mathscr{H}}\int_{\mathscr{Y}} \left\|u - u_{\mathsf{post}}^{\boldsymbol{y}}\right\|_{\mathscr{H}}^2 \pi_{\mathsf{like}}(\boldsymbol{y}|u)\, d\boldsymbol{y}\, \mu_{\mathsf{pr}}(du),$$

which is also known as the Bayes risk of the estimator $u_{\mathsf{post}}^{\boldsymbol{y}}$, corresponding to a quadratic loss function [5, 3]. The following result extends the well known result regarding the connection between the Bayes risk of the MAP estimator and the trace of the posterior covariance, for a Bayesian linear inverse problem, to the infinite-dimensional Hilbert space setting.

Theorem 2. *Let μ_{pr} be a centered Gaussian prior measure on $\mathscr{H}$, and $\mu_{post}^{\boldsymbol{y}} = \mathcal{N}(u_{post}^{\boldsymbol{y}}, \mathcal{C}_{post})$ be the posterior measure for a Bayesian linear inverse problem with additive Gaussian noise model as described in Section 3. Then, $\mathsf{E}_{\mu_{pr}}\left\{\mathrm{MSE}(u_{post}^{\boldsymbol{y}}; u)\right\} = \mathrm{tr}(\mathcal{C}_{post})$.*

Proof. By Lemma 4,

$$\mathsf{E}_{\mu_{\mathsf{pr}}}\left\{\mathrm{MSE}(u_{\mathsf{post}}^{\boldsymbol{y}}; u)\right\} = \int_{\mathscr{H}} \|(\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}} - I)u\|_{\mathscr{H}}^2 \; \mu_{\mathsf{pr}}(du) + \mathrm{tr}(\mathcal{C}_{\mathsf{post}}^2\mathcal{H}_{\mathsf{m}}), \qquad (22)$$

and since $(\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}} - I)u \sim \mathcal{N}(0, (\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}} - I)\mathcal{C}_{\mathsf{pr}}(\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}} - I)^*) =: \mu$,

$$\int_{\mathscr{H}} \|(\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}} - I)u\|_{\mathscr{H}}^2 \; \mu_{\mathsf{pr}}(du) = \int_{\mathscr{H}} \|\xi\|_{\mathscr{H}}^2 \; \mu(d\xi) = \mathrm{tr}((\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}} - I)\mathcal{C}_{\mathsf{pr}}(\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}} - I)^*).$$

We proceed as follows,

$$\begin{aligned}\mathrm{tr}((\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}} - I)\mathcal{C}_{\mathsf{pr}}(\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}} - I)^*) \;&=\; \mathrm{tr}((\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}} - I)^*(\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}} - I)\mathcal{C}_{\mathsf{pr}}) \\ &= \mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}^2\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}) - \mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}\mathcal{C}_{\mathsf{pr}}) - \mathrm{tr}(\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}) + \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}).\end{aligned}$$

Let us first consider the last two terms; recalling that $\mathcal{C}_{\mathsf{post}} = \mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{S}\mathcal{C}_{\mathsf{pr}}^{1/2}$, we have

$$\begin{aligned}-\mathrm{tr}(\mathcal{C}_{\mathsf{post}}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}) + \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}) &= -\mathrm{tr}(\mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{S}\mathcal{C}_{\mathsf{pr}}^{1/2}\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}) + \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}) \\ &= -\mathrm{tr}(\mathcal{C}_{\mathsf{pr}}\mathcal{S}\tilde{\mathcal{H}}_{\mathsf{m}}) + \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}) \\ &= \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}\mathcal{S}(\mathcal{S}^{-1} - \tilde{\mathcal{H}}_{\mathsf{m}})) = \mathrm{tr}(\mathcal{C}_{\mathsf{pr}}\mathcal{S}) = \mathrm{tr}(\mathcal{C}_{\mathsf{post}}).\end{aligned}$$

Thus, by (22), we have,

$$\mathsf{E}_{\mu_{\mathsf{pr}}}\left\{\mathrm{MSE}(u_{\mathsf{post}}^{\boldsymbol{y}}; u)\right\} = \mathrm{tr}(\mathcal{C}_{\mathsf{post}}^2\mathcal{H}_{\mathsf{m}}) + \mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}^2\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}) - \mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}\mathcal{C}_{\mathsf{pr}}) + \mathrm{tr}(\mathcal{C}_{\mathsf{post}}).$$

Hence, showing that the first three three sum to zero completes the proof. To this end, we note that $\mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}\mathcal{C}_{\mathsf{pr}}) = \mathrm{tr}(\tilde{\mathcal{H}}_{\mathsf{m}}\mathcal{S}\mathcal{C}_{\mathsf{pr}})$ and that

$$\begin{aligned}\mathrm{tr}(\mathcal{C}_{\mathsf{post}}^2\mathcal{H}_{\mathsf{m}}) + \mathrm{tr}(\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{post}}^2\mathcal{H}_{\mathsf{m}}\mathcal{C}_{\mathsf{pr}}) \;&=\; \mathrm{tr}(\tilde{\mathcal{H}}_{\mathsf{m}}\mathcal{S}\mathcal{C}_{\mathsf{pr}}\mathcal{S}) + \mathrm{tr}(\tilde{\mathcal{H}}_{\mathsf{m}}\mathcal{S}\mathcal{C}_{\mathsf{pr}}\mathcal{S}\tilde{\mathcal{H}}_{\mathsf{m}}) \\ &=\; \mathrm{tr}\big(\tilde{\mathcal{H}}_{\mathsf{m}}\mathcal{S}\mathcal{C}_{\mathsf{pr}}\mathcal{S}(I + \tilde{\mathcal{H}}_{\mathsf{m}})\big) = \mathrm{tr}(\tilde{\mathcal{H}}_{\mathsf{m}}\mathcal{S}\mathcal{C}_{\mathsf{pr}}). \quad \square\end{aligned}$$

Appendix A. Proof of Lemma 1

Let $\{e_i\}_1^\infty$ be a complete orthonormal set in $\mathscr{H}$, and denote by Π_n the orthogonal projection of $\mathscr{H}$ onto Span$\{e_1, \ldots, e_n\}$; that is, for $x \in \mathscr{H}$, $\Pi_n(x) = \sum_{i=1}^n \langle e_i, x\rangle_{\mathscr{H}} e_i$. First note that,

$$\int_{\mathscr{H}} \langle \mathcal{A}x, x\rangle_{\mathscr{H}} = \int_{\mathscr{H}} \langle \mathcal{A}(x-m), x-m\rangle_{\mathscr{H}}\ \mu(dx) + \int_{\mathscr{H}} \langle \mathcal{A}x, m\rangle_{\mathscr{H}}\ \mu(dx) + \int_{\mathscr{H}} \langle \mathcal{A}m, x\rangle_{\mathscr{H}}\ \mu(dx) - \langle \mathcal{A}m, m\rangle_{\mathscr{H}}\,.$$

Now by the definition of the mean of the measure, the last three terms sum to $\langle \mathcal{A}m, m\rangle_{\mathscr{H}}$. Thus, the rest of the proof consists of showing $\int_{\mathscr{H}} \langle \mathcal{A}(x-m), x-m\rangle_{\mathscr{H}}\ \mu(dx) = \mathrm{tr}(\mathcal{A}\mathcal{Q})$. Note that for every $x \in \mathscr{H}$, $x - m = \lim_{n\to\infty} \Pi_n(x-m)$, and thus,

$$\lim_{n\to\infty} \langle \mathcal{A}\Pi_n(x-m), \Pi_n(x-m)\rangle_{\mathscr{H}} = \langle \mathcal{A}(x-m), x-m\rangle_{\mathscr{H}}\,.$$

Moreover, we note that, $|\langle \mathcal{A}\Pi_n(x-m), \Pi_n(x-m)\rangle_{\mathscr{H}}| \le \|A\|\,\|x-m\|^2_{\mathscr{H}}$, and since $\mathcal{Q}$ is trace-class, the measure μ has a bounded second moment; hence, $\int_{\mathscr{H}} \|x-m\|^2_{\mathscr{H}}\ \mu(dx) < \infty$. Therefore, we can apply Lebesgue-Dominated Convergence Theorem to get,

$$\lim_{n\to\infty} \int_{\mathscr{H}} \langle \mathcal{A}\Pi_n(x-m), \Pi_n(x-m)\rangle_{\mathscr{H}}\ \mu(dx) = \int_{\mathscr{H}} \langle \mathcal{A}(x-m), x-m\rangle_{\mathscr{H}}\ \mu(dx). \tag{23}$$

Next, let $\{e_i\}_1^\infty$ be the (complete) set of eigenvectors of $\mathcal{Q}$ with corresponding (real) eigenvalues $\{\lambda_i\}_1^\infty$. We know that $\mathcal{A}\mathcal{Q}$ is trace-class with,

$$\mathrm{tr}(\mathcal{A}\mathcal{Q}) = \sum_i \langle \mathcal{A}\mathcal{Q}e_i, e_i\rangle_{\mathscr{H}} = \sum_i \lambda_i \langle \mathcal{A}e_i, e_i\rangle_{\mathscr{H}}\,. \tag{24}$$

Also, note,

$$\begin{aligned}
&\int_{\mathscr{H}} \langle \mathcal{A}\Pi_n(x-m), \Pi_n(x-m)\rangle_{\mathscr{H}}\ \mu(dx)\\
&\quad= \sum_{i,j=1}^n \int_{\mathscr{H}} \langle \mathcal{A}e_i, e_j\rangle_{\mathscr{H}} \langle x-m, e_i\rangle_{\mathscr{H}} \langle x-m, e_j\rangle_{\mathscr{H}}\, \mu(dx)\\
&\quad= \sum_{i,j=1}^n \langle \mathcal{A}e_i, e_j\rangle_{\mathscr{H}} \int_{\mathscr{H}} \langle x-m, e_i\rangle_{\mathscr{H}} \langle x-m, e_j\rangle_{\mathscr{H}}\, \mu(dx)\\
&\quad= \sum_{i,j=1}^n \langle \mathcal{A}e_i, e_j\rangle_{\mathscr{H}} \langle \mathcal{Q}e_i, e_j\rangle_{\mathscr{H}} = \sum_{i=1}^n \lambda_i \langle \mathcal{A}e_i, e_i\rangle_{\mathscr{H}}\,.
\end{aligned}$$

Therefore, combining this last result with (23) and (24), we get

$$\begin{aligned}
\int_{\mathscr{H}} \langle \mathcal{A}(x-m), x-m\rangle_{\mathscr{H}}\ \mu(dx) &= \lim_{n\to\infty} \int_{\mathscr{H}} \langle \mathcal{A}\Pi_n(x-m), \Pi_n(x-m)\rangle_{\mathscr{H}}\ \mu(dx)\\
&= \lim_{n\to\infty} \sum_{i=1}^n \lambda_i \langle \mathcal{A}e_i, e_i\rangle_{\mathscr{H}} = \mathrm{tr}(\mathcal{A}\mathcal{Q}).\ \square
\end{aligned}$$

References

[1] A. Alexanderian, N. Petra, G. Stadler, and O. Ghattas. A-optimal design of experiments for infinite-dimensional Bayesian linear inverse problems with regularized ℓ_0-sparsification. *SIAM Journal on Scientific Computing*, 2014. to appear.

[2] A. C. Atkinson and A. N. Donev. *Optimum Experimental Designs*. Oxford, 1992.

[3] J. O. Berger. *Statistical decision theory and Bayesian analysis*. Springer, 1985.

[4] T. Bui-Thanh, O. Ghattas, J. Martin, and G. Stadler. A computational framework for infinite-dimensional Bayesian inverse problems part i: The linearized case, with application to global seismic inversion. *SIAM Journal on Scientific Computing*, 35(6):A2494–A2523, 2013.
[5] B. P. Carlin and T. A. Louis. Bayes and empirical bayes methods for data analysis. *Statistics and Computing*, 7(2):153–154, 1997.
[6] K. Chaloner and I. Verdinelli. Bayesian experimental design: A review. *Statistical Science*, 10(3):273–304, 1995.
[7] J. B. Conway. *A course in operator theory*. American Mathematical Soc., 2000.
[8] G. Da Prato. *An introduction to infinite-dimensional analysis*. Springer, 2006.
[9] G. Da Prato and J. Zabczyk. *Second order partial differential equations in Hilbert spaces*, volume 293. Cambridge University Press, 2002.
[10] M. Dashti, K. J. Law, A. M. Stuart, and J. Voss. MAP estimators and their consistency in Bayesian nonparametric inverse problems. *Inverse Problems*, 29, 2013.
[11] H. P. Flath, L. C. Wilcox, V. Akçelik, J. Hill, B. van Bloemen Waanders, and O. Ghattas. Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial hessian approximations. *SIAM Journal on Scientific Computing*, 33(1), 2011.
[12] S. Kullback and R. A. Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 03 1951.
[13] F. Pukelsheim. *Optimal design of experiments*, volume 50. siam, 2006.
[14] M. Reed and B. Simon. *Methods of Modern Mathematical Physics: Vol.: 1.: Functional Analysis*. Academic press, 1972.
[15] B. Simon. Notes on infinite determinants of Hilbert space operators. *Advances in Mathematics*, 24:244–273, 1977.
[16] A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numerica*, 19:451–559, 2010.
[17] D. Uciński. *Optimal measurement methods for distributed parameter system identification*. CRC Press, Boca Raton, 2005.

Institute for Computational Engineering and Sciences, The University of Texas at Austin
E-mail address: alen@ices.utexas.edu

Mathematics Department, United States Naval Academy
E-mail address: gloor@usna.edu

Institute for Computational Engineering and Sciences, The University of Texas at Austin
E-mail address: omar@ices.utexas.edu